Practical Accounting in Content-Centric Networking (extended version)

Cesar Ghali Gene Tsudik Christopher A. Wood
University of California, Irvine
{cghali, gene.tsudik, woodc1}@uci.edu

Edmund Yeh
Northeastern University
eyeh@ece.neu.edu

arXiv:1510.01852v1 [cs.NI] 7 Oct 2015

***Abstract*—Content-Centric Networking (CCN) is a new class of network architectures designed to address some key limitations of the current IP-based Internet. One of its main features is in-network content caching, which allows requests for content to be served by routers. Despite improved bandwidth utilization and lower latency for popular content retrieval, in-network content caching offers producers no means of collecting information about content that is requested and later served from network caches. Such information is often needed for accounting purposes. In this paper, we design some secure accounting schemes that vary in the degree of consumer, router, and producer involvement. Next, we identify and analyze performance and security tradeoffs, and show that specific per-consumer accounting is impossible in the presence of router caches and without application-specific support. We then recommend accounting strategies that entail a few simple requirements for CCN architectures. Finally, our experimental results show that forms of *native and secure* CCN accounting are both more viable and practical than application-specific approaches with little modification to the existing architecture and protocol.**

I. Introduction

The original Internet was designed in the late 1970s with the main purpose of providing end-to-end communication. It allowed thousands of users to remotely access scarce computing resources from terminals. Since then, the number of Internet users has grown exponentially. They use a wide variety of applications, many of which involve some form of content distribution. This shift in usage exposed some design limitations of the current IP-based Internet and motivated the exploration of new networking architectures.

Content-Centric Networking (CCN) is an approach to internetworking exemplified by two well-known research efforts: the CCNx project at the Palo Alto Research Center [1] and Named-Data Networking (NDN) [2]. In IP-based networking, a user requests *content* by addressing the host at which it is stored. Conversely, in CCN, *content* is assigned a unique name and is addressed directly. Any entity can become a content *producer* as long as it can show that is authorized for a certain part of the global content namespace. A user, called a consumer, requests *content* by issuing an *interest* carrying the former's name. Such interests can be satisfied by any entity (host or router) that either creates or caches the requested content. The *content* follows, in reverse, the exact path of the preceding interest towards the consumer. Any intervening routers on this path may cache the content to satisfy future interests for the same content.

These in-network caches facilitate efficient content distribution. This important feature helps reduce end-to-end latency and lower bandwidth consumption when requesting popular content. Since an interest can be satisfied by a copy of the requested content found in a router's cache, it might not reach the producer. Consequently, a producer might only receive a small fraction of all interests for a given piece of content. At the same time, the number, sources, and timing of interests represent important information that could be used for accounting by the producer. Even if the timing and the number of interests were somehow communicated to the producer, interest sources would remain unknown since CCN lacks consumer information, e.g., source addresses, in interests.

Furthermore, router cache space will likely be treated as a valuable (and even premium) resource as CCN is deployed in the real world. Thus, a mechanism is needed for reporting cache hits to content producers or router owners, thereby informing them about content usage. To be viable, such a mechanism must only incur minimal bandwidth, computation, and storage overheads. Finally, to prevent attacks such as false cache usage reporting, it also must be be secure. In this paper, we design a lightweight secure accounting mechanism, applicable, to both CCNx and NDN[1]. Our intended contributions are three-fold:

- Identification and motivation for features needed for CCN accounting and for security thereof.
- The first comprehensive technique for content and cache usage accounting, with varying levels of consumer, router, and producer involvement.
- Analysis of performance and security tradeoffs.

In the rest of this paper, we use the term CCN to refer to both CCNx and NDN.

Organization. Section II overviews CCN. Next, Section III discusses desired features for content accounting. Security requirements are addressed in Section IV. Performance of the proposed approach is assessed in Section VI. The paper concludes with a summary of related work in Section VII.

II. CCN Overview

This section gives an overview of CCN. Given familiarity with CCN, it can be skipped without loss of continuity. Note that all details are presented in the context of the CCNx architecture and protocol. Specifics such as packet formats, message fields, and routing decisions have subtle differences in NDN. However, with minor changes to the protocol and packet formats, our description also applies to NDN.

[1]Support for NDN requires minor packet format and protocol changes.

Unlike IP which focuses on end-points of communication and their names/addresses, CCN stresses content by making it named, addressable, and routable [1]. A content name is represented with a labeled content identifier (LCI) string, which is similar to a standard URI. For example, a BBC news home-page for April 15, 2015 might be named `lci:/uk/bbc/news/2015APR15/index.htm`.

CCN communication follows the *pull* model whereby content is delivered to consumers only upon explicit request. There are two basic types of packets (messages) in CCN: interest and content object messages[2]. A consumer requests content by issuing an *interest* message. If an entity can "satisfy" a given interest, it returns the corresponding *content object*. Name matching in CCN is exact, i.e., an interest for `lci:/snapchat/Bob/video-749.avi` can only be satisfied by a content object named `lci:/snapchat/Bob/video-749.avi`.

CCN interest messages include, at a minimum, the name of the requested content. They may also carry a payload that enables consumers to push data to producers along with the content request [3].[3] CCN content objects include a number of fields. In the context of this paper, we are only interested in the following four:

- `Name` – A URI-compliant sequence of name components.
- `Payload` – The actual data of the content object.
- `Validation` – A composite of validation algorithm information (e.g., the signature algorithm used, its parameters, and a link to the public verification key), and validation payload (e.g., the signature). For simplicity, we use the term "signature" to refer to this field.
- `ExpiryTime` – Producer-recommended time for the content objects to be cached.

There are three types of CCN entities or roles:[4]

- *Consumer* – an entity that issues an interest for content.
- *Producer* – an entity that produces and publishes content.
- *Router (Forwarder)* – an entity that routes interest packets and forwards corresponding content packets.

Each CCN entity maintains three structures [4]:

- *Content Store* (CS) – a cache used for content caching and retrieval. Cache size is determined by local resource availability. Each router may unilaterally determine whether to cache content and for how long. From here on, we use the terms *CS* and *cache* interchangeably.
- *Forwarding Interest Base* (FIB) – a table of name prefixes and corresponding outgoing interfaces. The FIB is used to route interests based on longest-prefix-match of the name.
- *Pending Interest Table* (PIT) – a table of outstanding (pending) interests and a set of corresponding incoming interfaces.[5]

[2]We use the terms *content* and content object interchangeably.

[3]Currently, NDN interest messages do not provide an arbitrary-length payload and therefore cannot support some of the proposed accounting types, as will become clear later in the paper. However, if in the future the NDN interest format includes a field similar to the CCNx payload, our accounting approach will become applicable.

[4]A physical entity (a host, in today's parlance) can be both consumer and producer of content.

[5]A given PIT entry might indicate multiple incoming interfaces reflecting the possibility of multiple interests arriving at, or near, the same time.

Upon receiving an interest with the name N (i.e., for the content with name N), a router first checks its cache for existence of a local copy of content with the same name.[6] If a local copy is not found in the cache, and there are no pending interests for N, the router forwards the interest to the next hop(s) according to its FIB and forwarding strategy. For each forwarded interest, a router creates a new PIT entry with state information, including the name and the interface on which it arrived. Moreover, if an interest with the name N arrives while there is already an entry for the same name in the PIT, the router collapses the present interest and only stores the interface on which it was received. When content is returned, the router forwards it out on all recorded incoming interfaces and flushes the corresponding PIT entry. Since no additional information is needed to deliver content, an interest does not carry a source address.

III. Accounting in CCN

As mentioned earlier, router caches present a major challenge for accounting in CCN. In particular, if interests are satisfied by caches, **how can a content producer collect information about the popularity of (or demand for) its content**? In this section, we discuss design elements for accounting in CCN. For the time being, we do not emphasize security and privacy considerations. Specifically, assuming benign (well-behaved) consumers, routers, and producers, our initial goal is to determine the minimal functionality needed by all CCN entities to facilitate *correct* accounting. In doing so, we consider three types of accounting information:

- *Individual:* This type of information is tied directly to a specific consumer. An example might be the number of times a particular consumer requested a particular content. It provides linkability between consumers and content they obtain. Moreover, it requires revealing consumer identities, at least to the producer.
- *Distinct:* This type of information is functionally equivalent to the individual accounting information with the exception that consumer identities are not revealed. Instead, a randomly generated nonce is added to the each interest to enable a producer to distinguish between separate interests. More details are explained in the following sections.
- *Aggregate:* This type of information represents an aggregate over a set of consumers. For example, it might include the number of times a particular piece of content was requested from a specific geographic location or an ISP. Aggregate information enables some degree of consumer privacy.

We believe that these three types are sufficiently representative of any accounting information needed in any real-world CCN application and focus on them in the remainder of this work. Also, we claim that accounting should not be mandatory for all content. Some producers might not care about the popularity of *any* of their content, while others might need accounting information only for *some* of their content. We refer to content for which producers desire such information as *accountable* content.

[6]This is why CCN lacks any notion of a *destination address*; since content can be served by any CCN entity.

Lastly, one important design dimension is whether accounting information is reported in real time (online) or not (offline). In the latter case, a network management protocol can be envisaged whereby an AS-level accounting server periodically collects cache hit logs from its routers and reports the results to individual producers or AS-housing producers. This kind of accounting seems viable. However, it involves a potentially significant delay in notifying producers about the demand, or lack thereof, for their content. This might be unacceptable for content for which real time demand information is needed. Intuitively, real-time accounting is a more difficult problem to solve. We therefore limit the rest of this paper to real-time accounting.

A. Counting Cache Hits vs. Content Requests

Another variable in supporting CCN accounting is exactly what is being counted: instances of cache hits, or instances of requested content being served to the consumer? A cache hit occurs when a router or another entity finds the requested content in its cache. By *another entity* we mean either a producer or a repository that keeps a copy of the content. We assume that accounting for cache hits is only relevant for routers, i.e., network elements.[7] An instance of content being served occurs when a cache hit takes place **and** the content is actually delivered to a *single* consumer.

It might seem that these two types of events are the same, i.e., a content is served once for every cache hit, and vice-versa. However, this is not the case in CCN. Whenever a router receives an interest, it may choose to multicast (forward) it out on multiple interfaces. This behavior is officially allowed since a router's FIB can express multiple next hops for a given name prefix. One practical reason for allowing it that it facilitates quick(er) fetching of content. However, it also complicates accounting. Consider the following scenario:

> Suppose that a consumer requests content CO and the issued interest is received by router R_1. The latter then forwards the interest for CO to two upstream routers R_2 and R_3 based on its FIB. Both R_2 and R_3 have CO in their respective caches and each replies to R_1 with its cached version. Assuming that R_2's copy of CO is the first to reach R_1, the latter forwards CO downstream and flushes the appropriate PIT entry. When R_3's copy of CO arrives, R_1 discards it since it does not refer to a current existing PIT entry. If both R_2 and R_3 inform CO's producer P about a cache hit on CO, P would incorrectly assume that CO was requested twice. Even though, technically, CO is served twice by two distinct routers, there was only one requesting consumer which received only one copy of CO.

According to this scenario, the count of cache hits can exceed the content request count. This problem occurs because there is no way to distinguish among different interests issued for the same content.[8] In other words, if consumers Cr_1 and Cr_2 issue interests for COvat different times, their interests would be identical. Moreover, even if CO is not cached, i.e., interests for it reach P, and if Cr_1 and Cr_2 issue interests for CO at roughly the same time, P would be unable to distinguish between this case – when two consumers ask for CO – and the case in the scenario above – when one consumer asking for CO and R_1 decides to multicast the interest upstream. Note that the number of cache hits is equal to two in both cases, but the number of content objects served is two and one, respectively.

The reason for supporting both types of accounting is quite intuitive: a producer might need to know the exact demand for its content, whether on an aggregate or individual basis. Separately, a producer might need to know which routers experience cache hits for its content. The latter could be used to reconcile billing the producer for cache usage.

Finally, even though accounting for cache hits and content requests is not the same thing, we naturally would like to use the same mechanism as much as possible to provide both. Therefore, for the rest of the paper and unless otherwise mentioned, we use the term *accounting* to refer to both accounting for cache hits and content requests.

B. Accounting via Content Encryption

One intuitive accounting approach is to use encrypted content.[9] Suppose that producers encrypt all accountable content, and the decryption keys – which, in CCN, are represented as content objects with well-defined names – are configured not to be cached, i.e., by setting their `ExpiryTime` to 0. Even if consumer interests requesting such content are satisfied from in-network caches, the former must separately issue interests requesting the decryption key(s). Such interests bypass in-network caches and reach the producers, thereby enabling per-request accounting.

With content encryption, the desired type of accounting dictates how interests requesting keys should be generated. In the case of individual accounting, consumers must include some kind of consumer-specific data in the interests when keys are requested. Such data allows producers to link these interests to specific consumers. However, if only aggregate accounting is required, interests requesting keys do not have to carry any consumer-specific data. As mentioned above, such interests need to have some kind of a nonce to enable the producer to distinguish between the case of receiving two interests from two different consumers, or receiving two interests sent from a single consumer and were multicasted by a router in the network.

Accounting via content encryption has two primary advantages: (1) it is transparent to the network layer, and (2) it does not require any new features and message types. However, despite its apparent simplicity, it is not efficient. All accountable content objects need to be encrypted and keys need to be requested and distributed separately. Thus, content is obtained by issuing two interests – one for the content and one for the key(s) – thus incurring at least two round-trips to the producer.

[7] Accounting for cache hits in repositories or at producers themselves is out of scope.

[8] NDN interests carry a random nonce used for interest loop detection, which can be helpful in this distinction. However, CCNx interests do not carry nonces.

[9] This is also a form of access control.

We believe that an ideal accounting mechanism should efficiently work for all accountable content. That is, it should not require a consumer to issue more than a single interest for accountable content. Also, general content accounting should be distinct from content access control.

C. Accounting via Push Interests

The accounting approach proposed in this paper is based on real-time reporting. The key element is a new message type that we call a *push* interest, denoted as *pInt*. Its main purpose is to inform the producer that its content has been requested and a cache hit occurred. Structurally, a *pInt* carries a name similar to a regular interest. However, the most important distinguishing feature of a *pInt*, is that, unlike a regular interest, it does not leave behind any state in routers. Specifically, a *pInt* referencing content CO is forwarded by each router until it reaches its corresponding producer P, and no information about that *pInt* is retained by any intervening router. A router forwards a *pInt* just like it forwards a regular interest with the exception that *pInt* messages are not multicasted. This restriction is necessary to prevent producers from receiving duplicate copies of the same *pInt*.

Besides this forwarding change, the behavior of CCN routers is slightly modified to support *pInt* generation. There are two cases when a router generates a *pInt*:

1) Whenever a regular interest with name N is satisfied from its cache, a router generates a *pInt* with the same name N and forwards it upstream towards the producer.
2) When a router receives a content object corresponding to a PIT entry, it forwards that message on all downstream interfaces listed in said PIT entry. However, before flushing that entry, a router generates a *pInt* that aggregates all collapsed interests. (These aggregation details are discussed in Section III-D.) Note that *collapsed* refers only to those interests that were not originally forwarded upstream. This is because the one forwarded upstream presumably already (1) reached the producer, or (2) triggered its own *pInt* via case 1 above.[10]

In order for a producer to inform routers about what content requires accounting, we also introduce a new `ACCT` flag in the content header, which reflects one the following three values:

1) `NONE`: the producer requests no accounting information for this content.
2) `AGGREGATE`: the producer requests aggregate accounting information for this content.
3) `DISTINCT`: the producer requests distinct accounting information for this content.
4) `INDIVIDUAL`: the producer requires individual interest-level accounting for this content.

Whenever a cache hit occurs, routers behave the same for cases 2, 3, and 4. The only difference is when a content arrives and a router has a number of previously collapsed interests for that content. In case 2, a router generates a *pInt* with the count of collapsed interests for a given content. In cases 3 and 4, a router reports the *actual* interests, which can optionally be bundled into a single *pInt*.

As hinted in the above discussion, we aim to support aggregate, distinct, and individual accounting types, and to report instances of cache hits as well as instances of content requests. Distinct and individual accounting for cache hits and content requests (as well as aggregate for cache hits) is possible and indeed attained by the proposed technique. Unfortunately, as will become clear below, supporting accurate distinct and aggregate accounting for *content requests* is quite challenging. However, if some kind of consumer-specific data is provided in interest `Payload` fields, probabilistically accurate distinct and aggregate accounting for content requests can be achieved.

D. pInt Format and Features

We now describe the *pInt* message format and discuss the purpose of its fields. Structurally, this message is nearly isomorphic to CCNx 1.0 interests [5], and includes the following necessary fields:

- `Name`: copied entirely from the `Name` field in the interest (or PIT entry) that triggers a *pInt*.
- `Type`: flag indicating whether this *pInt* is for aggregate, distinct, or individual accounting.
- `Origin`: identifies the router that generates the *pInt*, e.g., the router's prefix (if available) or public key digest.
- `Count`: set to 1 in the case of a cache hit, or the number of interfaces minus one on which the content object was forwarded downstream if interest collapsing occurred.
- `Cdata`: a random nonce or consumer-specific data used by producer for different purposes based on the accounting type required (i.e., distinct or individual). If `Count`>1, this is a sequence of `Count` of consumer-specific data culled from corresponding interests. Such data can be carried in the interest `Payload` field.

The semantics of the `Cdata` field depend on the type of required accounting information. As stated above, aggregate accounting for cache hits does not require `Cdata` to be present. In the following we discuss consumer-specific data requirements for other accounting types.

Aggregate accounting for content requests: The problem in this type of accounting is that producers do not have the means to distinguish between the cases where received interests (or *pInt* messages) with the same name are multicasted by routers or generated by several distinct consumers. However, if consumers include random nonces and timestamps as consumer-specific data, this distinction can be achieved.

Individual Accounting for cache hits and content requests: In this case, `Cdata` needs to reflect the identify of consumers issuing interests. This value can take a variety of forms:

1) Consumer public keys or their digests. Note that this form reveals consumer identities to all network entities – not only producers.
2) Group public keys or their digests. A group can be an organization, autonomous system (AS), or a geographical region. In this case, the group identify is revealed rather than that of the individual consumer.

[10] For example, suppose that a router receives an interest for content CO on interfaces: 2, 3, 5, and 6. Regularly, only the first one (arriving on interface 2) is forwarded, say, on interface 9. Interests on 3, 5 and 6 arrive later and are collapsed in the same PIT entry. Now, when CO arrives back on interface 9, it is forwarded out on all 4 incoming interfaces. However, the router generates a *pInt* that reflects only the last 3 interfaces: 3, 5 and 6.

3) Unique consumer identifiers (i.e., pseudonyms). Although this does not violate consumer anonymity, such identifiers need to be assigned to consumers by producers or a trusted third party before any interests are issued. Note that this form of `Cdata` allows interest linkability.[11]
4) Consumer identity (using any of the three previous forms) with nonces and timestamps. This form of consumer-specific data allows producers to know which consumers request what content, as well as how many times such requests are made.

Distinct Accounting for cache hits and content requests: For this accounting type, `Cdata` needs to reflect the uniqueness of interests. This can be achieved if consumers include a nonce and timestamp in the `Cdata` field. The format of the nonce is application-specific and can range from a random number to the hash of the content name and the timestamp.
Note that the same knowledge provided to the producer in the distinct accounting case can also be attained using aggregate accounting type if `Cdata` reflects the uniqueness of interests. However, we keep the distinction between these two types for ease of classification.

Each of the above forms impose different overhead on consumers and producers. However, router overhead is only very slightly affected if accounting is done with *pInt* messages.[12] This is because routers simply populate the `Cdata` field of generated *pInt* messages using information contained in the `Payload` field of the corresponding interests, regardless of how consumer-specific data is generated. In other words, *routers are oblivious to the accounting type used*.

Also, note that the choice of which form to use is an application-specific issue. We do not mandate a specific technique.

E. Correctness

We define correctness of an accounting technique as follows.

Definition 1. *An accounting technique is* correct *if it accurately reports cache hit and content request information to the producer, assuming that all participants faithfully follow the rules (i.e., exhibit no malicious behavior) and there are no transmission errors, packet loss, or node failures that affect accounting-relevant traffic.*

We also define probabilistically correct accounting technique as follows.

Definition 2. *An accounting technique is* probabilistically correct *if it is correct with a negligible probability of error, i.e., reporting inaccurate or false information.*

We now informally prove correctness of each of the proposed accounting techniques: individual, distinct, and aggregate in both cache hits and content requests cases.

Cache Hits: A router R generates *pInt* messages for producer P for every cache hit on accountable content objects. Since all routers (including R) on the path to P forward *pInt* messages according to their FIB entries, such messages are guaranteed to be received by P.[13] This provides accurate cache hit counts for content objects in question. This argument holds for all three types of accounting. The only difference is that `Cdata` fields of *pInt* must contain appropriate consumer-specific data in some accounting types.

Content Requests: Recall that *pInt* messages are also generated to report collapsed interests. Such interests reflect content requests in the network. Although *pInt* messages provide *correct* individual, distinct, and aggregate accounting for cache hits, it can only provide *probabilistically correct* accounting for content requests. As stated above, consumers can include nonces and timestamps in the `Payload` fields of their interests. If producers post-process received *pInt* messages, this `Payload` information allows them to distinguish between the cases where received interests (or *pInt* messages) with the same name are multicast by routers or generated by several distinct consumers. However, since nonces are random strings of bits, collisions might occur with a probability negligible in their bit length.

F. Performance

As previously mentioned, our goal is to design a lightweight mechanism for secure accounting in CCN. In this section we discuss involvement of consumers, routers and producers.

Consumers: The degree of consumer overhead varies depending on the accounting type required. In the case of aggregate accounting for cache hits, *pInt*-based accounting is transparent to the consumers. However, when aggregate accounting for content requests is required, distinct or individual accounting techniques must be used. In this case, consumers must include a random nonce or consumer-specific data in all issued interests as described in Section III-C.[14] The overhead for consumers depends on how this data is generated. For instance, consumer-specific data resilient to some types of attacks can increase this overhead – see Section IV-B for details.

Routers: Regardless of the accounting type used, router involvement is always required when *pInt* messages are used.[15] For all accountable content, routers must generate *pInt* messages whenever cache hits occur or collapsed interests are satisfied. However, we argue that routers overhead is minimal.

To enable efficient throughput, a router must be able to quickly construct *pInt* messages in the fast path. All *pInt* field values can be copied from their corresponding counterparts in interest and content headers. For instance, the `Name` and `Type` values can be copied from the corresponding content header fields `Name` and `ACCT`. The `Cdata` field can be populated by copying `Payload` field values from corresponding interests. Moreover, the router identifier in the `Origin` field can be pre-computed (if needed) and stored in memory.

Algorithm 1 illustrates a procedure to generate *pInt* messages in routers. It is triggered whenever a router receives a

[11] Interest linkability is defined as the ability of an eavesdropper (observer or adversary) to reveal that fact whether two captured interests are issued by the same consumer.

[12] Content encryption-based accounting is completely transparent to routers.

[13] Recall that we assume there are no transmission errors or packet loss events.

[14] Human consumers are not involved, only their applications.

[15] Content encryption-based accounting is transparent to routers.

TABLE I. CCN ENTITIES OVERHEAD IN ACCOUNTING. N/A ALSO INDICATES THAT NO INVOLVEMENT IS REQUIRED.

	Consumers		**Routers**		**Producers**	
Accounting Type	**Aggregate**	**Individual**	**Aggregate**	**Individual**	**Aggregate**	**Individual**
Cache hits	N/A	consumer-specific data generation	*pInt* messagesgeneration		$\mathcal{O}(\lvert\mathbb{C}_P\rvert)$ counters	$\mathcal{O}(\lvert\mathbb{U}_P\rvert \cdot \lvert\mathbb{C}_P\rvert)$ counters & consumer-specific data sanitation
Content requests	nonces/timestamps					

Algorithm 1 *pInt*-Generation

```
1: Input: CO[N, ACCT], Int[N, PL], R_id
2: pInt.Name := CO.N
3: pInt.Type := CO.ACCT
4: pInt.Origin := R_id
5: if CO from local cache then
6:     pInt.Cdata := Int.PL
7:     pInt.Count := 1
8: else
9:     e = FindPITEntry(CO.N)
10:    for each i in e/{Int} do
11:        pInt.Cdata := pInt.Cdata || i.PL
12:        pInt.Count := pInt.Count + 1
13:    end for
14: end if
15: Forward pInt according to the FIB
```

solicited content object or locates a copy in its local cache. The algorithm takes as input the received content object CO with name N and accounting flag `ACCT`, its corresponding interest Int with name N and payload PL, and the router R_{id}. If CO is not in local cache, the router (1) copies consumer-specific data from all interests in the corresponding PIT entry into `Cdata`, and (2) sets `Count` value accordingly.

To obtain better better bandwidth utilization, routers can generate a single *pInt* message for a specific content CO in a pre-defined time window. In this case, routers report *batched* cache hits by (1) including the actual cache hit counters in the *pInt*'s `Count` field, and (2) listing all consumer-specific data from the corresponding interests (if any) in the *pInt*'s `Cdata` field.

Producer: Similar to routers, producers requesting accounting information are always involved and the overhead varies depending on the type of accounting required. We describe these differences below.

- Aggregate: the overhead is minimal in this case. Producers only need to maintain a cache hit counter for every accountable content they publish. The counters maintained at producer P are in order $\mathcal{O}(|\mathbb{C}_P|)$, where $\mathbb{C}_P$ is the set of accountable content published by P.
- Individual: producers have to maintain a counter for every {consumer, accountable content} pair. The size of the counters table is in order $\mathcal{O}(|\mathbb{U}_P| \cdot |\mathbb{C}_P|)$, where $\mathbb{U}_P$ is the set of consumers requesting accountable content from P. The producer might also store the timestamp at which requests for content arrived. This, obviously, storage requirements. Moreover, to provide *correct* individual accounting for content requests, producers must post-process all received *pInt* messages in order to detect duplicates caused by interest multicasting by routers.
- Distinct: in this case, producers have to maintain a counter for every accountable content object as well as information about all requests for these objects. Such information could be the timestamp at which requests are received, as discussed in Section III-C. The size of the stored information is in order $\mathcal{O}(|\mathbb{R}_P|+|\mathbb{C}_P|)$, where $\mathbb{R}_P$ is the set of all requests received by P for its published accountable content.

Table I summarizes each CCN entity overhead in the proposed accounting mechanisms.

IV. SECURITY CONSIDERATIONS

Thus far, we assumed that all entities involved in accounting are benign, i.e., act honestly and correctly. However, this assumption is clearly unrealistic in practice. In this section, we propose security requirements to achieve secure accounting in CCN. We also show that some attacks cannot be prevented or even detected without additional cryptographic overhead. Moreover, we stress that implementing secure accounting incurs a trade-off between security and overhead on consumers and producers. Fortunately, routers are unaffected by such amendments.

A. Adversary Model

The anticipated adversary Adv is a malicious router generating *pInt* messages for non-existing interests, when individual accounting is required. In other words, Adv tries to inflate individual accounting information for both cache hits and content requests. We also assume that consumers behave honestly: if a consumer needs to provide consumer-specific data or a random nonce, it does so correctly. We consider dishonest consumers later in Section V.

For completeness, we identify the following additional attacks and show why we exclude them from the discussion.

- A router that (1) does not generate *pInt* messages when necessary, or (2) generates *pInt* messages without forwarding content downstream. Both cases can be reduced to packet loss events. We do not address these attacks since the router's malicious behavior cannot be detected.
- A consumer that continuously generates interests to inflate accounting information. If aggregate or distinct accounting is required, the producer will not be able to detect such malicious behavior.
  On the other hand, if individual accounting is required, consumer-specific data can be used to detect continuous requests. However, this scenario can be reduced to an Interest Flooding (IF) attack [6], [7], which is outside the scope of this paper. A similar argument applies for distinct accounting information.
- An external attacker that controls network links, that is capable of eavesdropping on, dropping, or replaying packets, including *pInt* messages. This attack is irrelevant if links are encrypted, which is a realistic assumption for

adjacent routers. In most cases, consumers and producers are connected to edge routers using link-layer encryption, e.g., EAP [8].

Again, we do not consider an adversary that tries to inflate aggregate or distinct accounting information. As previously argued, this cannot be prevented deterministically due to the likely usage of multicast forwarding strategies. Therefore, we only consider security of correct individual accounting information. More formally, we define a secure accounting technique as follows.

Definition 3. *An accounting technique is* secure *with respect to Adv if it is correct and all Adv malicious behavior can be detected.*

B. Mitigating Forgeries and Replay Attacks

Section III-C mentioned several options for generating consumer-specific data. However, in order to prevent inflation attacks, such data must be unforgeable and resistant to replay attacks. We define *secure* consumer-specific data as follows.

Definition 4. *Consumer specific data is* secure *if it can be authenticated by at least the producer, and is neither forgeable nor subject to replay attacks.*

Providing replay resiliency can be accomplished if consumer-specific data carries a nonce r and a timestamp t. If producers receive consumer-specific data with duplicate r in the same time window to which t belongs, such data is discarded. Using these values, secure consumer-specific data Sec-CrSD takes the following form.

$$\mathsf{Sec\text{-}CrSD} = \Big[\mathsf{CrSD}||r||t, \quad \mathsf{f}_k\left(\mathsf{CrSD}||r||t||Int.N\right)\Big] \tag{1}$$

where CrSD is consumer-specific data formed according to Section III-D[16], $\mathsf{f}_k(\cdot)$ is a function that provide unforgeability, and k is a secret key (either symmetric or private key.) Note that interest name in the $\mathsf{f}_k(\cdot)$ computation is used to bind Sec-CrSD to the interest to which it is appended. This prevents Adv from using the same Sec-CrSD for generating multiple *pInt* messages with different names. The function $\mathsf{f}_k(\cdot)$ can be a Hash Message Authentication Code (HMAC) [9] (if consumers share secrets with producers) or a digital signature generation function. Each method has well-known advantages and drawbacks. Symmetric HMAC-based functions are generally much less costly than digital signatures. However, the former requires *a priori* key distribution. On the other hand, using digital signatures involves in-line distribution and on-line verification of signers' public keys. This method of generating Sec-CrSD incurs a lot of computational and storage overhead. In addition to signature verification, producers are required to maintain a list of all received nonces in the acceptable current time window for each accountable content.

At this point, it becomes clear that unforgeability and replay resiliency cannot be achieved unless *secure* consumer-specific data is used. This is not possible if aggregate or distinct accounting is required, where consumer-specific data is not provided. One way to solve the problem is to include Sec-CrSD in all interests regardless of the accounting type required. This is impractical as it introduces unnecessary overhead for both consumers and producers.

[16]Note that CrSD and `Cdata` are different. The former is generated by consumers and assigned to `Payload` field of interest, while the latter is a field in a *pInt* and may contain none or many CrSD-s.

C. Preserving Consumer Anonymity

In order to provide *secure* individual accounting in the presence of Adv, consumer-specific data should be generated securely. Digital signatures, by nature, reveal the consumer's identity, and HMAC tags allow separate interests to be linked together since a key identifier must be included in the interest to properly verify each HMAC. To this end, we develop a technique for generating Sec-CrSD anonymous to all network entities, except producers.

We begin with the notion of consumer-specific data indistinguishability, which is necessary to maintain anonymity among an arbitrary set of consumers.

Definition 5. *Let Cr^a and Cr^b be two consumers generating two consumer-specific data values CrSD^a and CrSD^b in two different interests for the same content object CO. Let $\mathcal{E}$ be any eavesdropper (except the producer publishing CO) and not directly connected to either Cr^a or Cr^b. Let the event of $\mathcal{E}$ successfully revealing the source of CrSD^a and CrSD^b be denoted as $\mathcal{E}_{\text{rev}}\left(\mathsf{CrSD}^a, \mathsf{CrSD}^b\right) = 1$. We claim that these two interests are indistinguishable if the probability of $\mathcal{E}_{\text{rev}}\left(\mathsf{CrSD}^a, \mathsf{CrSD}^b\right) = 1$ is no better than a random guess. That is,*

$$\Pr\left[\mathcal{E}_{\text{rev}}\left(\mathsf{CrSD}^a, \mathsf{CrSD}^b\right) = 1\right] \leq \frac{1}{2} + \epsilon(n),$$

for any negligible function ϵ and security parameter n.

Moreover, we assume that consumers know the producer's public key pk before requesting its content; see Section V for justification. Let A-CrSD denote an anonymous consumer-specific data of the form:

$$\mathsf{A\text{-}CrSD} = \mathsf{Enc}_{pk}\left(\mathsf{Sec\text{-}CrSD}\right) \tag{2}$$

where $\mathsf{Enc}_{pk}(\cdot)$ is a public key encryption function using pk, and Sec-CrSD is formed according to Equation 1.

To prevent $\mathcal{E}$ from learning that two interests are generated by the same consumer, A-CrSD values in the two interests should be indistinguishable. This can only be achieved if $\mathsf{Enc}_{pk}(\cdot)$ is a CPA-secure public key encryption scheme, i.e., secure against Chosen Plaintext Attacks [10]. In some encryption schemes, this is done by mixing in a random number (nonce) with the plaintext before encryption.

Theorem 1. *Assume a CPA-secure public key encryption scheme $\mathsf{Enc}_{pk}(\cdot)$. An* A-CrSD *composed according to Equation 2 guarantees* indistinguishable *consumer-specific data with a negligible probability of nonce collision.*

Proof: The proof of consumer-specific data indistinguishability follows from the proof of CPA-secure public key encryption scheme [10]. We only prove that A-CrSD generation guarantees negligible probability of nonce collision.

We assume individual accounting, and that f is the frequency in which consumers send interests to a specific producer during a specific time window w, where each consumer generates appropriate A-CrSD values in the interests. Let the number of interests sent be $s = f \times w$. We claim that the probability of any two A-CrSD-s in the set $\{\text{A-CrSD}_1, \ldots, \text{A-CrSD}_s\}$ being derived from colliding nonces is negligible in N, the length of the nonce in bits. Let this collision event be denoted as $\text{Col}(\text{A-CrSD}_i, \text{A-CrSD}_j)$ for $i \neq j$ and $1 \leq i, j \leq s$. The probability of this event occurring can be calculated according to the birthday paradox as follows.

$$\begin{aligned}
&\Pr\Big[\text{Col}(\text{A-CrSD}_i, \text{A-CrSD}_j) = 1 \ ; \ i \neq j \ , \ 1 \leq i, j \leq r\Big] \\
&= 1 - \left(\frac{2^N}{2^N} \times \frac{2^N - 1}{2^N} \times \cdots \times \frac{2^N - s + 1}{2^N}\right) \\
&= 1 - \frac{2^N!}{(2^N)^s \, (2^N - s)!} \\
&= 1 - \frac{s! \cdot \binom{2^N}{s}}{(2^N)^s} \qquad (3)
\end{aligned}$$

Note that Equation 3 assumes that $s < 2^N$; otherwise, the collision probability is equal to 1 according to the Pigeonhole Principle. ■

We note that public-key operations to generate and verify A-CrSD values may be overly expensive for some consumers. In such scenarios, a potential optimization is to use symmetric encryption. Specifically, A-CrSD values can be computed as $\text{enc}_k(\text{Sec-CrSD})$, where $\text{enc}_k(\cdot)$ is a CPA-secure symmetric encryption function and k is a shared secret between the consumer and the producer. This, however, requires additional operations for creating and managing shared secrets. In order for producers to determine which shared secret to use for decrypting received A-CrSD-s, consumers should add a form of cleartext identifier, e.g., shared key tag. Although this violates interests indistinguishability, it does indeed preserve consumer anonymity.

V. Individual Accounting in Practice

In this section we describe some challenges related to transparently collecting individual accounting information. We then outline recommendations for applications wishing to implement accounting in CCN.

A. Individual Accounting Challenges

Thus far, we made several assumptions about consumers as far as individual accounting information conveyed to producers:

1) Consumers know *what* accounting information is needed for a desired content object.
2) Consumers know the producer's public key pk used to encrypt Sec-CrSD sections.
3) Consumers behave honestly, i.e., for content that requires individual accounting information in the CrSD, consumers supply *correct* required information.

The first assumption seems to be particularly problematic, especially if a given consumer has no prior relationship with a specific producer. However, there are at least two simple ways for consumers to learn what a producer expects in an interest.

First, recall that adherence to sound trust management at the network layer [11] requires the consumer to know, when requesting content (issuing an interest), the public key of the producer. This, in turn, means that the consumer prefetched the producer's public key. It is easy to extend the producer's public key certificate to include desired accounting requirements for constructing correct interests for that producer and the namespace included in the certificate.

One natural alternative is for consumer to "blindly" issue a trial interest for some random content in the namespace of a given producer. This interest might not abide by the producer's accounting rules. However, the producer could then reply with its public key certificate that would include the producer's accounting requirements.

Without performing one of the aforementioned techniques, a consumer cannot be expected to provide specific information in an interest since it does not have any expectation of the form of this information. We consider possible consumer behavior for each accounting type below.

- Aggregate: For this type of accounting information, the CrSD fields are empty and *pInt* messages only contain a count. Thus, consumers are not required to have any a priori information when constructing their original interest.
- Distinct: CrSD contains random nonces for each interest issued, and *pInt* messages carry collections of these nonces from consumers to producers along with a total count. The nonce does not require any application-specific input from the consumer to generate; it is simply a random string. Thus, in this case, consumers are also not required to have any a priori information.
- Individual: For this type of accounting information, CrSD contains a very specific piece of consumer-specific data (identifier) for each interest, and *pInt* messages propagate these values to producers. To be useful for individual accounting, producers must be able to utilize the provided information in order to identify each consumer. Furthermore, it will likely be the case that different producers require different identification information, both in its content and representation. Thus, consumers cannot be expected to know *which* type of consumer-specific data to provide in an interest without having been told beforehand.

It is clear that individual accounting information necessitates some initial interaction or registration phase, wherein consumers are given the interest CrSD requirements and also the public key used when generating Sec-CrSD or A-CrSD values in interests. Note that the issue of public key identification is analogous to the problem of not knowing *what* consumer-specific data to use for a given interest.

This interaction or registration step covers assumptions (1) and (2) above, but it does not address assumption (3). Namely, even if a consumer has all of the information at their disposal needed to construct a valid interest for a given content object, what happens if they maliciously choose to use the *wrong* information? Such an adversary can easily obtain data from router caches without having to provide the correct accounting

information to the producer, thereby effectively bypassing the accounting mechanism.

The core problem is that routers have no means to determine if the information contained in interest CrSD fields is *correct*[17] for a given content object. To be able to verify this information routers would have to possess (or be told) some piece of information for each cached content that requires individual accounting. Only then can routers verify CrSD field values before replying with a cached content. Not only is this unreasonable for routers, it also means that anonymous consumer-specific data, as described in Section IV-C, is no longer feasible. Since routers must be able to verify the CrSD for each interest associated, the consumer must necessarily reveal *some* information about their identity.

Therefore, we claim that assumption (3) is not realistic at the network-layer in the presence of caches and dishonest consumers. This means that individual accounting must be handled at the application layer. The proof of this claim can be argued from the above discussion.

B. Recommendations

Given the previous discussion, we present some recommendations for collecting accounting information in CCN. First, if individual accounting information is needed, producers must simply set all content cache time to zero (0). This will force all interests to be routed to the producer without being satisfied from an in-network cache. If an interest for content that requires individual accounting is received and the required accounting information is missing, producers should reply with a Negative Acknowledgment (NACK) [12] indicating consumer-specific data requirements to obtain said content. The consumer can then re-issue an interest with the correct information. If consumers go through a preliminary registration step, this accounting requirement information can be obtained once and then used for all subsequent interests, thereby removing the need for an additional round-trip. Observe that since producers process all interests before responding with content, they can determine if a given interest for individual accountable content is valid and thus detect behavior by malicious dishonest consumer.

For aggregate and distinct accountable content, consumers should always generate a random nonce and include it in CrSD. If a router caches the content and its `ACCT` flag is `AGGREGATE`, then CrSD values can simply be dropped when *pInt* messages are generated. Otherwise, if the router caches the content and its `ACCT` flag is `DISTINCT`, the nonce is copied into the generated *pInt*. This is a simple modification to the router *pInt* generation procedure described in Algorithm 1, and induces no significant overhead for consumers or routers.

This simple policy can be extended to *all interests*. Since consumers are not generally expected to know what type of accounting information is required for content, they can blindly generate a nonce for each interest they issue. CCN routers will then correctly propagate these nonces in *pInt* messages to the producer according to the rule above. As previously noted, NDN already supports default nonce generation in interests (but for the purposes of interest loop detection). The CCNx protocol needs to be updated to include this requirement.

[17]This means that the provided information is the one required by the producer.

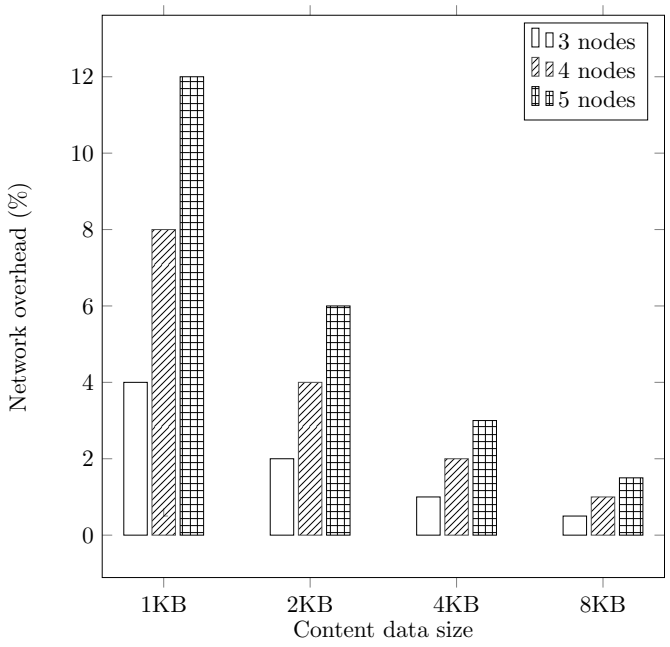

Fig. 1. Network overhead imposed by forwarding *pInt* messages.

VI. Analysis and Experimental Assessment

In Section III, we proposed two fundamental techniques for propagating accounting information to producers: encryption-based and *pInt*-based solutions. The former technique is beneficial in that it is entirely transparent to the routers. Encryption-based accounting, which is a form of access control, is an application-layer concern, and therefore the routers do not require any modification to support the scheme. Conversely, accounting based on *pInt* messages requires routers to execute the *pInt*-Generation procedure upon every cache hit to generate *pInt* messages, and also forward *pInt* messages towards producers using the same data plane logic as normal interest messages.

Consider a scenario in which we have k consumers $Cr_1, \ldots, Cr_k$ and a single producer P. Let $Cr_i, R_1, \ldots, R_l, P$ be a consumer-to-producer path traversed by interests issued by consumer Cr_i for the accountable content object CO. Let R_c, $1 \leq c \leq l$ be the router at which CO is cached. Furthermore, let p_l be the number of messages traversing the $R_1 - R_c$ path in one direction, and let p_r be the number of messages traversing the $R_c - P$ path in one direction. Finally, let γ be the number of content requests issued by all consumers Cr_i, $i = 1, \ldots, k$, along the $R_1 - P$ path. Recall that encryption-based accounting requires consumers to issue at least two (2) interests to access accountable content: one interest is issued for the content itself, and then at least one more is issued to request the corresponding decryption keys. The former interest will traverse the $R_1 - R_c$ path, whereas all decryption key interests will traverse the full $R_1 - P$ path. Thus, in this case, $p_l = 4\gamma$ and $p_r = 2\gamma$. Conversely, in the *pInt*-based approach, a single interest is issued for CO on the $Cr - R_c$ path, and then a *pInt* is generated at R_c and forwarded along the $R_c - P$

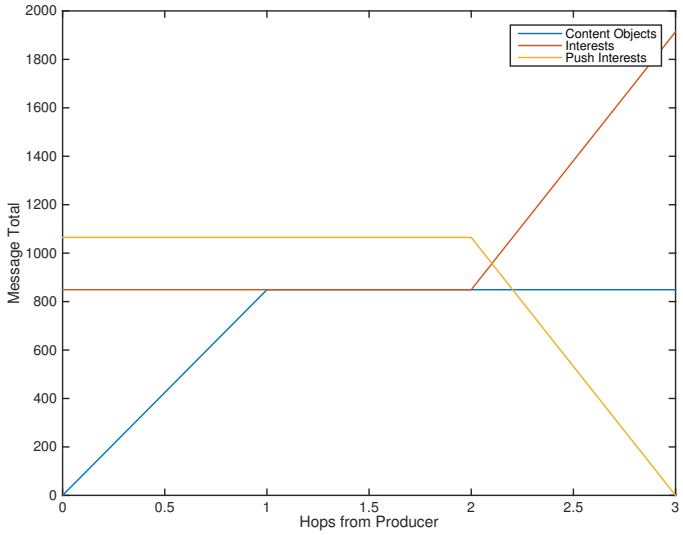

(a) Path topology with 5 nodes (1 consumer, 3 routers, and 1 producer), $A = 500$ and $M = 1000$.

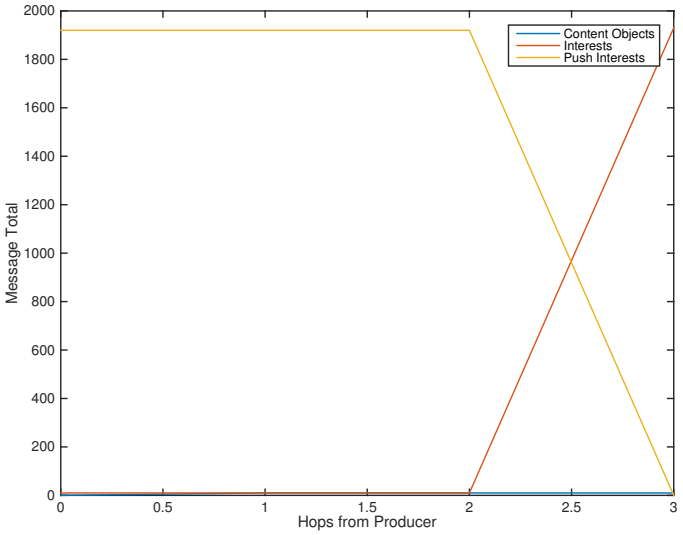

(b) Path topology with 5 nodes (1 consumer, 3 routers, and 1 producer), $A = 500$ and $M = 10$.

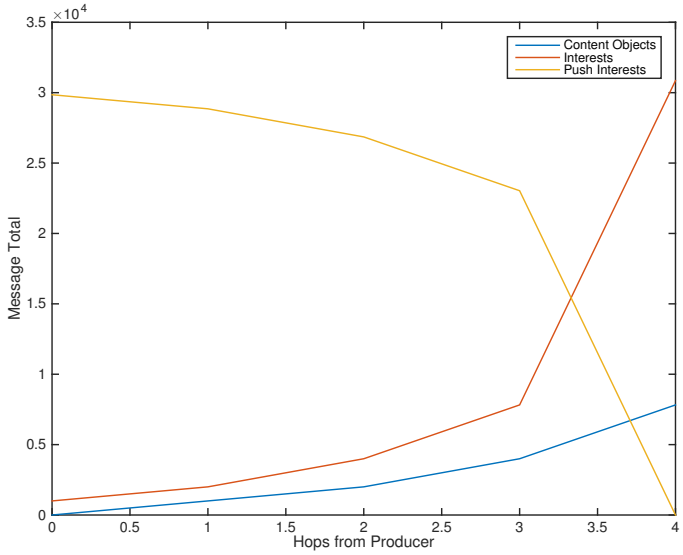

(c) Binary tree of height 5 (32 consumers, 30 routers, and 1 producer), $A = 500$ and $M = 1000$.

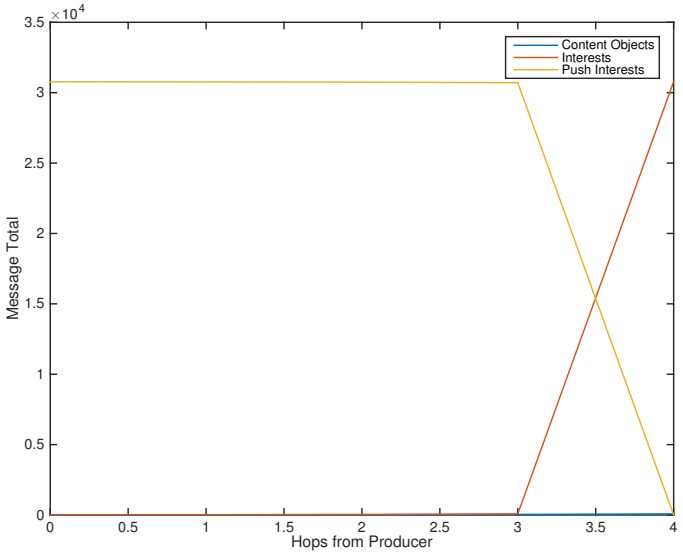

(d) Binary tree of height 5 (32 consumers, 30 routers, and 1 producer), $A = 500$ and $M = 10$.

Fig. 2. *pInt*-based accounting overhead in networks with path and tree topologies.

path. Therefore, in this case, $p_l = 2\gamma$ and $p_r = \gamma$. Note that the case where $R_c = P$ is identical to the scenario where there are no in-network caches, in which case there would be no *pInt* messages generated. This case performs worse than the *pInt*-based variant since $p_l = 2\gamma$ and $p_r = 2\gamma$ (there is a single RTT from the Cr_i to P for CO).

Notice that the differences in p_l are due to the fact that, unlike interests, *pInt* messages have no response from the producer. In fact, the *network overhead*, in terms of the number of messages, of the encryption-based accounting solution is exactly twice that of the *pInt*-based solution, and the network overhead of the cacheless variant (which, again, invalidates the need for *pInt* messages and accounting information) is more than the *pInt*-based solution as well. Furthermore, producers, and consumers incur additional overhead since encryption and decryption must be performed, respectively, in order to consume content. Therefore, in our initial experimental assessment, we limit our focus to the *pInt*-based accounting solution, since we feel that it is (a) a more efficient technique and (b) also proportional to the overhead incurred by network entities in the encryption-based scenario. In our experimental analysis, we assume that adding the *pInt* generation procedure is a constant time operation for routers. We also assume that any shared-secret management protocols, such as those that might be used to establish or acquire a shared HMAC key, are done offline and are therefore not part of the real-time or online communication.

If interests are satisfied from the cache of R_c, all upstream routers on the consumer-to-producer path incur the overhead of forwarding *pInt* messages to the producer. Figure 1 shows

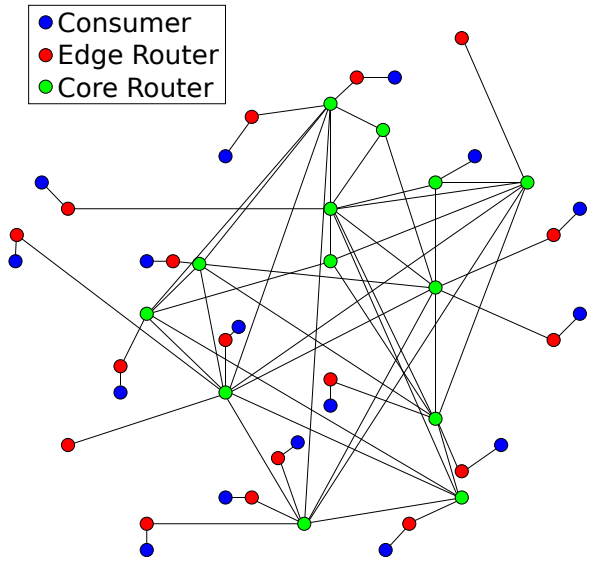

Fig. 3. DFN topology - each edge router above is connected to a group of NDN consumers.

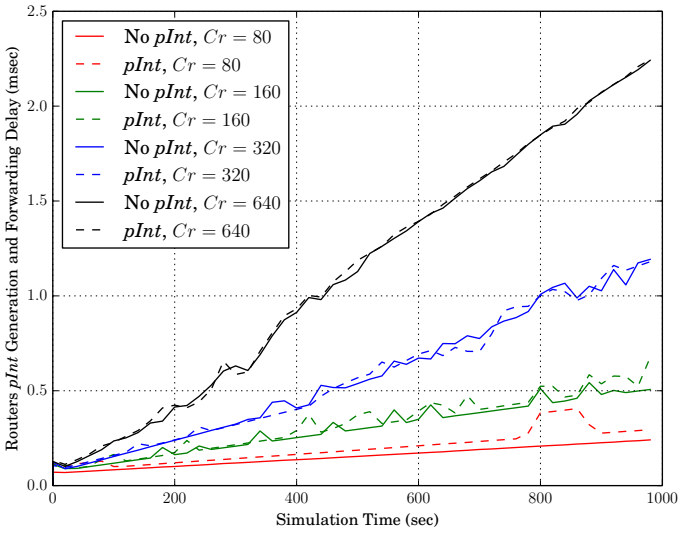

Fig. 4. DFN topology - *pInt* messages generation overhead at routers.

this overhead as a function of corresponding content size and the number of links between the router generating *pInt* and the producer. It is calculated as the the ratio of the extra bytes (due to the forwarded *pInt* messages) traversing each link over the size of the corresponding content object. The x-axis represents content data size, without including the header. We calculated the overhead in three line topologies containing 3, 4, and 5 nodes consisting of 2, 3, and 4 links respectively. The first node is the consumer Cr and the last node is the producer P. For the purpose of this exercise, we assume the following:

- Any content requested by Cr can always be satisfied by the first hop (consumer facing) router's (R) cache, i.e., cache hit rate is 100%. This accounts for the highest network overhead since *pInt* must traverse all links except the first one connecting Cr with R.
- Each router's FIB is pre-configured to forward all interests and *pInt* messages towards P.

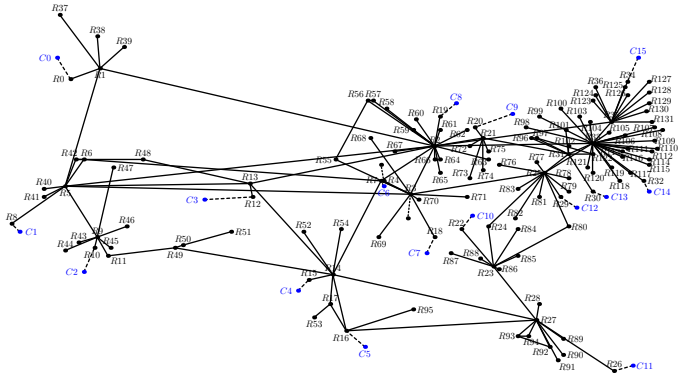

Fig. 5. AT&T topology - each edge router above is connected to a group NDN consumers.

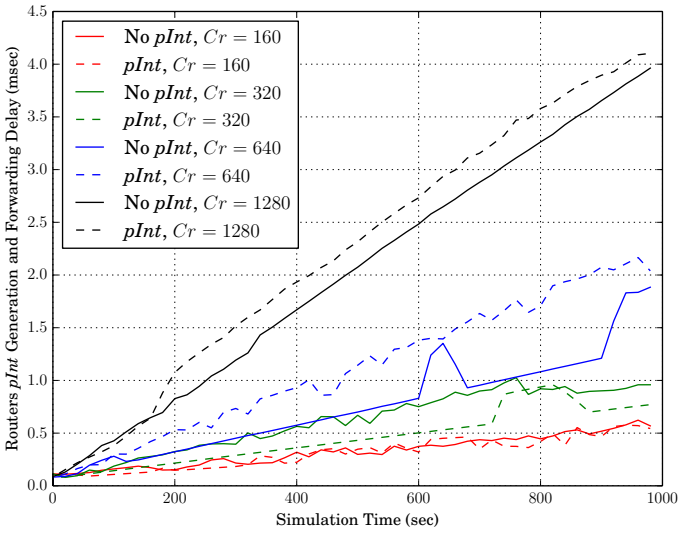

Fig. 6. AT&T topology - *pInt* messages generation overhead at routers.

- Interest, *pInt*, and content object headers only contain a name of length 40B.

The results from this experiment show that, as the size of content objects grow, the bandwidth overhead induced by *pInt* messages decreases. This overhead would increase if more complex topologies were considered, e.g., k-ary trees rooted at P. However, it would see the same decline as the content object size increased.

A. *Message Count Overhead*

To further understand the performance impact of *pInt*-based accounting in different topologies, we studied the amount of overhead incurred by each entity – consumer, router, and producer – in the network as a function of the distance from the producer. To do this, we implemented a custom discrete-time event-driven simulator that models a variety of CCN network topologies: paths and binary trees. Consumers are configured to issue interests for a single producer at a Poisson rate with parameter A. The names of each interest are uniformly sampled from a pool of M names. Each router invokes the *pInt*-Generation procedure upon every cache hit. We place no restriction on cache sizes, since the set of possible content objects is small enough to fit within any reasonable sized

cache. Specifically, producers respond to interests with content objects with a fixed payload size of 1MB.

We argue that the number of messages is indicative of the bandwidth overhead induced by *pInt* messages. This is because the size of *pInt* messages will be proportional to the size of interests, even with *secure* consumer-specific data. Therefore, it is not the size of these messages that is important – it is the sheer quantity which are processed by network elements. Thus, we are concerned with the number of messages that propagate from routers to the producer.

We study this overhead in networks with path and tree topologies. For simplicity, we restrict our analysis to 5-hop paths and binary trees of height 5. By varying A and M, we show how many messages of each type are processed by each entity as a function of the distance from the producer. Figure 2 illustrates the obtained results. Clearly, as M decreases, the likelihood of cache hits increases. This results in a clear increase in *pInt* processing at each upstream entity from the cache hit location. For example, in the cases where $M = 10$, approximately 99% of all messages processed by routers upstream of cache locations were *pInt* messages.

The interest request rate is highly dependent on the type of application. High request rates for popular content, which is likely to be cached, will lead to a proportionally high number of *pInt* messages propagating upstream to the producer. If interests are issued for unpopular or uncached content, then approximately the same number of interests will be propagated upstream. In other words, from the perspective of the producer, the sum of the received interests and *pInt* messages will be equal the total number of content requests from all consumers: **the producer overhead is linear in the number of content requests.** The difference in these two cases is that *pInt* messages are typically smaller in size than their interest counterparts.

B. Router Overhead Assessment

To measure router overhead incurred by generating and forwarding *pInt* messages we extended ndnSIM [13], a simplified implementation of NDN architecture as a NS-3 [14] module for simulation purposes, to support *pInt* messages. With this modified architecture, we two sets of experiments using two different topologies:

- The DFN network, Deutsches ForschungsNetz (German Research Network) [15], [16]: a German network developed for research and education purposes which consists of several connected routers positioned in different areas of the country, as shown in Figure 3. The network consists of a total of 30 routers. Blue dots in the figure represent group of consumers connected to edge routers (red dots), while green dots represent core network routers.
- The AT&T backbone network [7]: shown in Figure 5, this network consists of more than 130 routers, and each logical consumer in the figure represents multiple physical consumers connected to an edge router.

In all of our experiments, consumers issue interests at a rate of 10 interest per second for the same content with the name `/prefix/A/00`. To capture the worst case scenario, wherein the maximum number of *pInt* messages are generated, we (1) disable interest collapsing, and (2) set the `ExpiryTime` of the request content to be equal to the simulation time, ensuring that this content is cached at routers throughout the whole duration of the simulation. This forces routers to generate a single *pInt* for every cache hit, resulting in the maximum amount of *pInt* messages that can be possibly generated.

We measure the overhead required by routers to generate and forward *pInt* messages as compared the the case where *pInt* messages are not generated. Figure 4 shows the router overhead in the DFN topology parameterized by the number of consumers connected to edge routers (80, 160, 320, and 640). We observed that even with 640 consumers in the network, the overhead of an average router when generating *pInt* messages is negligible. Similarly, Figure 6 shows the overhead of generating and forwarding *pInt* messages by routers in the AT&T topology. In the case with 1280 consumers, routers experience a 15% additional overhead while generating *pInt* messages, which we consider to be negligible and a difference that can be recovered with better routing hardware.

VII. Related Work

Network-layer accounting in CCN and related interest-based ICN architectures remains an open topic in the literature [17]. However, certain economic aspects, such as *how* to set and enforce prices, has been widely discussed [18], [19]. These results imply an application-layer strategy whereby payment (not usage) information is willingly sent on behalf of the consumer. This conflicts with the approach advocated by Agyapong et al. [20], wherein only ISP-related entities (i.e., not consumers or producers) are involved in payment coordination. [20] considers producer payment as an application layer concern. Our accounting techniques facilitates a blend between these two schemes wherein usage and payment information are sent *autonomously* on behalf of the network-layer for consumers (end-hosts) and routers. ISP entities and producers are informed of usage information for billing purposes, and can follow up with payment collection at a later point in time.

Patané et al. [21] study a similar problem in the context of IP-compatible architectures. Specifically, they focus on ones with dedicated router caches like Content Distribution Networks (CDNs) and transit networks that chauffeur traffic between different ISP provider networks. Payment policies proposed in [21] are identical, though. All parties pay for the resources which were used to deliver their content. Patané et al. also neglect to discuss means by which this payment and usage information can be propagated. Similar to [19], Kocak et al. [22] discuss methods where content providers can coordinate price information and contracts between ISPs. Kocak et al. also opt for an open, unfederated approach, which fits with our model of autonomous accounting information propagation.

Another important element of this work is the generation of secure consumer-specific data in *pInt* messages. There is rich literature of packet-level authentication in the IP-based Internet, much of which is contained in [23]. However, techniques such as digital signatures and symmetric-key MACs require some possibly unrealistic assumptions, such as shared keys amongst all pairs of routers and trusted third parties for key generation and management. Using improved public-key cryptographic algorithms based on elliptic curves can

help improve the signature scheme efficiency [24], as with DNSCurve [25]. However, the sheer volume of interests in CCN and related ICNs will very likely be substantially larger than DNS queries in IP networks, leading to only relatively modest improvements in performance.

VIII. Conclusion

This paper represents the first attempt to address accounting in CCN. It presented a simple and lightweight accounting technique and showed how to enhance it with security without significant burden to consumers, producers, and routers. We analyzed performance of the proposed technique and demonstrated that secure accounting is both possible and practical in CCN.

References

[1] "Content centric networking (CCNx) project," http://www.ccnx.org.

[2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *CoNEXT*, 2009.

[3] "CCNx 1.0 protocol specications roadmap," http://www.ccnx.org/113/ccn-publications-papers/ccn-1-0-protocol-roadmap-available/.

[4] "CCNx protocol: CCNx node model," http://www.ccnx.org/releases/latest/doc/technical/CCNxProtocol.html.

[5] "CCNx messages in TLV format," https://tools.ietf.org/html/draft-mosko-icnrg-ccnxmessages-01.

[6] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in named data networking," in *ICCCN*, 2013.

[7] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding ddos attacks in named data networking," in *LCN*, 2013.

[8] L. Blunk and J. Vollbrecht, "RFC 2284: PPP extensible authentication protocol (EAP)," 1998.

[9] H. Krawczyk, M. Bellare, and R. Canetti, "RFC 2104: HMAC: Keyed-hashing for message authentication," 1997.

[10] J. Katz and Y. Lindell, *Introduction to modern cryptography: principles and protocols*. CRC Press, 2007.

[11] C. Ghali, G. Tsudik, and E. Uzun, "Network-layer trust in named-data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 12–19, 2014.

[12] A. Compagno, M. Conti, C. Ghali, and G. Tsudik, "To nack or not to nack?" in *ICCCN*, 2015.

[13] A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnsim: NDN simulator for NS-3," *University of California, Los Angeles, Technical Report*, 2012.

[14] "Network simulator 3 (NS-3)," http://www.nsnam.org/.

[15] "DFN-Verein," http://www.dfn.de/.

[16] "DFN-Verein: DFN-NOC," http://www.dfn.de/dienstleistungen/dfninternet/noc/.

[17] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 2, pp. 1024–1049, 2014.

[18] A. Araldo, D. Rossi, and F. Martignon, "Design and evaluation of cost-aware information centric routers," in *ICN*, 2014.

[19] T.-M. Pham, S. Fdida, and P. Antoniadis, "Pricing in information-centric network interconnection," in *IFIP Networking Conference*, 2013.

[20] P. K. Agyapong and M. Sirbu, "Economic incentives in information-centric networking: Implications for protocol design and public policy," *IEEE Communications Magazine*, vol. 50, no. 12, pp. 18–26, 2012.

[21] R. Patané and J. Remond, "Economics of information-centric networks," *Internet Economics VIII*, p. 21, 2014.

[22] F. Kocak, G. Kesidis, T.-M. Pham, and S. Fdida, "The effect of caching on a model of content and access provider revenues in information-centric networks," in *SocialCom*, 2013.

[23] D. Lagutin, "Redesigning internet-the packet level authentication architecture," *Licentiates Thesis-Helsinki University of Technology*, 2008.

[24] D. B. Johnson and A. J. Menezes, "Elliptic curve DSA (ECDSA): an enhanced DSA," in *USENIX*, 1998.

[25] D. J. Bernstein, "DNSCurve: Usable security for DNS," 2009, http://dnscurve.org/.